Institute for Scientific Computing Research

Seminar Series Abstracts

(in reverse chronological order)

September 28, 2001

QMTest and QMTrack: Testing Code and Tracking Progress

Mark Mitchell
CodeSourcery, LLC

Email: mark@codesourcery.com

Abstract:

The QMTest and QMTrack tools are the result of a Software Carpentry (www.software-carpentry.com) development project. These portable, open source tools use an extensible, programmable design to allow customization by users.

An overview of the design of both tools, as well as a demonstration of QMTest, will be provided. CodeSourcery's plans for the future of these tools will also be discussed.

The audience will be encouraged to provide suggestions about features that would be of particular utility in their development projects.

Research web page: http://www.software-carpentry.com/

Institution web page: http://www.codesourcery.com/

September 27, 2001

Simulation of Compressible Turbulent Flows With Reaction

Sutanu Sarkar
University of California, San Diego

Email: sarkar@mechanics.ucsd.edu

Abstract:

Two outstanding problems in the area of compressible reactive flows will be discussed. The first problem concerns large eddy simulation where the strongly nonlinear dependence of the reaction rate term on temperature leads to a subgrid contribution in the resolved-scale equations, which must be modeled. We have developed a subgrid model by using the information available in the resolved scales along with additional physical input, namely, a model spectrum for the unresolved scales. Promising results have been obtained when evaluating this subgrid reaction rate model against a direct numerical simulation of a shear layer. The second problem occurs when a burn initiated in a NIF capsule encounters inhomogeneities of mixture fraction due to the introduction of inert shell material into DT mix by Rayleigh–Taylor or Richtmeyer–Meshkov instabilities. Preliminary results have been obtained in a simple model problem to identify the important parameters that control the modified burn propagation.

Speaker's web page: http://www-mae.ucsd.edu/RESEARCH/SARKAR/sarkar.html

Institution web page: http://www.ucsd.edu/

September 25, 2001

The Behavior of Converging Shock Waves of Spherical and Polyhedral Form

Don Schwendeman

Rensselaer Polytechnic Institute

Email: schwed@rpi.edu

Abstract:

In this talk the behavior of converging spherical shock waves will be discussed. An analysis of converging shocks, whose initial shape takes the form of regular polyhedra, will be presented within the framework of Whitham's theory of geometrical shock dynamics. The analysis of this problem is motivated by the earlier work on converging cylindrical shocks discussed in Schwendeman and Whitham (*Proc. R. Soc. Lond.*, A413, 297-311, 1987). In that paper, exact solutions were reported for converging polygonal shocks in which the initial shape reforms repeatedly as the shock contracts. For the polyhedral case, the analysis is performed both analytically and numerically for an equivalent problem involving shock propagation in a converging channel with triangular cross section.

It is found that a repeating sequence of shock surfaces composed of nearly planar pieces develops, although the initial planar surface does not reform, and that the increase in strength of the shock at each iterate in the sequence follows the same behavior as that for a converging spherical shock independent of the convergence angle of the channel. In this sense, the shocks are stable and the result is analogous to that found in the two-dimensional case discussed in the earlier paper. A numerical study of converging spherical shocks subject to smooth initial perturbations in strength shows a strong tendency to form surfaces composed of nearly planar pieces, suggesting that the stability result is fairly general.

Institution web page: http://www.rpi.edu/

September 14, 2001

Multilevel Conjugate Gradient Methods for Nonlinear Least-Squares Finite Element Computations

Gerhard Starke
University of Hannover

Email: starke@ifam.uni-hannover.de

Abstract:

Least-squares, mixed-finite element discretizations of nonlinear elliptic boundary value problems lead to nonlinear algebraic least-squares problems of large dimension. This talk describes the use of nonlinear conjugate gradient methods for these problems. The gradient of the functional and, consequently, the search directions are computed with respect to the natural norm of the underlying variational problem. This is crucial in order to achieve convergence rates that are independent of the refinement level of the triangulation. An appropriate multilevel inner product on the space where the variational problem is formulated allows for the efficient implementation of the method. We conclude with computational results for a sequence of discretizations based on adaptively refined triangulations for a variably saturated subsurface flow problem.

Speaker's web page: http://www.ifam.uni-hannover.de/~starke/

Institution web page: http://www.uni-hannover.de/

September 6, 2001

Introduction to Architecture and Performance of First European TeraFlops Supercomputer: The Hitachi SR8000

Gerhard Wellein
University of Erlangen-Nuernberg

Email: gerhard.wellein@rrze.uni-erlangen.de

Abstract:

Last year, the LRZ Munich released Europe's first TeraFlops computer for public user service, a Hitachi SR8000-F1. The Hitachi SR8000 series, which is based on RISC technology, was designed to incorporate the advantages of vector and scalar/RISC computing. This talk introduces the concept of pseudo-vector processing at the processor level as well as the flexible usage model, which supports both MPP-like and vector/SMP-like programming style. Performance characteristics for basic kernel loops and complex large-scale application codes will be discussed, including comparisons with other present-day MPP and vector supercomputers.

Research web page: http://www.hlrs.de/

September 6, 2001

Input/Output Scalability of Different Architectures

Preethy Vaidyanathan
University of California, Santa Cruz

Email: vaidyanathan1@llnl.gov

Abstract:

We characterize the I/O behavior of a computational biology application on Linux clusters with different file systems and on an IBM SP/2. This application plays a central role in the Human Genome Project. We show that locality is a very important factor affecting the performance of this application. We present the design of a user-level library for a new model of location-transparent storage to automatically redirect read accesses to the most appropriate location.

Institution web page: http://www.ucsc.edu/

August 31, 2001

Prefetching and Caching Strategies for Modern Memory Systems

Steven Reinhardt
University of Michigan

Email: stever@eecs.umich.edu

Abstract:

Microprocessor technology advances continue to provide larger on-chip caches and higher memory bandwidths, but off-chip memory latency is and will remain a significant bottleneck for system performance. Our research concerns how best to exploit growing capacity and bandwidth to counter the impact of latency.

We seek to exploit memory bandwidth by improving the effectiveness of prefetching. Traditional prefetching schemes risk degrading performance unless accuracies are high. Our prefetching framework effectively eliminates bandwidth contention by prioritizing requests, and reduces pollution effects via cache placement. Using this framework, aggressive prefetching of large memory regions boosts performance dramatically when applications have spatial locality.

We seek to exploit the capacity of on-chip caches more fully by managing their contents intelligently. Virtual-memory systems use full associativity and software-controlled replacement to combat the high latency of disk transfers; we are exploring the extent to which these ideas can be applied to large on-chip caches. We have developed a practical fully associative cache structure and a novel software replacement algorithm tailored for secondary caches. In tandem, they can achieve significant miss-rate reductions compared to conventional organizations.

Future work includes investigating the potential of application and operating-system involvement in secondary-cache management.

Speaker's web page: http://www.eecs.umich.edu/~stever/

Institution web page: http://www.umich.edu/

August 28, 2001

A Comparison of Factorization-free Eigensolvers with Application to Cavity Resonators

Peter Arbenz
Swiss Federal Institute of Technology

Email: arbenz@inf.ethz.ch

Abstract:

We investigate eigensolvers for the generalized eigenvalue problem $Ax = \lambda Mx$ with symmetric A and symmetric positive definite M that do not require the factorization of either A or M. We compare various variants of preconditioned Rayleigh quotient minimization and preconditioned Jacobi–Davidson algorithm by means of large-scale finite element discretizations originating from the design of the new RF cavity ring cyclotron installed at the Paul Scherrer Institute (PSI) in Villigen, Switzerland.

Institution web page: http://www.epfl.ch/

August 24, 2001

Interactive Visualization of Large Graphs and Networks

Tamara Munzner
Compaq Systems Research Center

Email: tamara.munzner@compaq.com

Abstract:

Many real-world domains can be represented as large node-link graphs: backbone Internet routers connect with 70,000 other hosts, mid-sized Web servers handle between 20,000 and 200,000 hyperlinked documents, and dictionaries contain millions of words defined in terms of each other. Computational manipulation of such large graphs is common, but previous tools for graph visualization have been limited to datasets of a few thousand nodes.

Visual depictions of graphs and networks are external representations that exploit human visual processing to reduce the cognitive load of many tasks that require understanding of global or local structure. We assert that the two key advantages of computer-based systems for information visualization over traditional paper-based visual exposition are interactivity and scalability. We also argue that designing visualization software by taking the characteristics of a target user's task domain into account leads to systems that are more effective and scale to larger datasets than previous work.

This talk contains a detailed analysis of three specialized systems for the interactive exploration of large graphs, relating the intended tasks to the spatial layout and visual encoding choices. We present two novel algorithms for specialized layout and drawing that use quite different visual metaphors. The H3 system for visualizing the hyperlink structures of Web sites scales to datasets of over 100,000 nodes by using a carefully chosen spanning tree as the layout backbone, 3D hyperbolic geometry for a Focus+Context view, and provides a fluid interactive experience through guaranteed frame rate drawing. The Constellation system features a highly specialized 2D layout intended to spatially encode domain-specific information for computational linguists checking the plausibility of a large semantic network created from dictionaries. The Planet Multicast system for displaying the tunnel topology of the Internet's multicast backbone provides a literal 3D geographic layout of arcs on a globe to help MBone maintainers find misconfigured long-distance tunnels.

Each of these three systems provides a very different view of the graph structure, and we evaluate their efficacy for the intended task. We generalize these findings in our analysis of the importance of interactivity and specialization for graph visualization systems that are effective and scalable.

Speaker's web page: http://graphics.stanford.edu/~munzner/

Institution web page: http://www.compaq.com/

August 22, 2001

3D Pore Structure/Fluid Distribution Measurement and Network Model Flow Studies

Brent Lindquist

State University of New York, Stony Brook

Email: lindquis@ams.sunysb.edu

Abstract:

X-ray computed microtomography (XCMT) has made it possible to create three-dimensional images of the pore space of rock at resolutions around 5 microns. A 1024^2 digitized image thus "captures" a cube of rock with a measurement of one-half cm on a side. The availability of such images has revitalized efforts to understand the connection between pore microstructure and core scale rock flow properties.

Our contribution to this effort has been the development of image analysis algorithms to extract 3D stochastic descriptions of pore geometry parameters. We present results for an analysis of a suite of Fontainbleau sandstones from 7.5 percent to 22 percent porosity. Parameters measured include distributions of pore coordination number, channel length, throat size, and pore volume, as well as nearest neighbor correlation distributions for throat-throat and throat-pore sizes.

This experimental analysis provides perfect input for network model flow simulations. The results obtained (average coordination numbers around 3.5, log-normal pore-size distributions, exponentially tailed throat size distributions) have profound impact on network models, which, in the absence of precise experimental measurements, have resorted to idealized geometrical networks. We present results performed in conjunction with our collaborators at Australian National University, demonstrating the profound effects that the realistic network models have on simulation results for residual oil saturations in invasion-percolation based models for two-phase flow.

Experimental techniques to image fluids in XCMT are beginning to improve. Because XCMT is non-invasive, time series data is possible. We discuss preliminary results on the intra-pore distributions of water and oil as a result of both water- and oil-driven hysteretic flows.

Speaker's web page: http://www.ams.sunysb.edu/~lindquis/lindquist.html

Institution web page: http://www.sunysb.edu/

August 21, 2001

Multidimensional Upwinding, Preconditioning, and Other Stuff from Michigan

Philip Roe
University of Michigan

Email: philroe@engin.umich.edu

Abstract:

This is a brief synopsis of current research in the University of Michigan Aerospace Engineering Department related to computational fluid dynamics. The faculty includes Bram van Leer, Ken Powell, Iain Boyd, and me. Topics include magnetohydrodynamics, rarified flow, elastic waves, and high-order schemes.

Methods that incorporate PDE information at a higher-than-one-dimensional level are discussed. Potential or realized benefits include:

1. Valid incompressible limits of compressible codes
2. Faster convergence to steady state
3. Solution-driven grids
4. Economically captured discontinuities
5. Preservation of div/curl constraints
6. Avoidance of computational anomalies

Institution web page: http://www.umich.edu/

August 13, 2001

Computational Study of Arterial Flow Disturbance Induced by Intravascular Stents

Levanto Schachter
University of California, Davis

Abstract:

A common interventional protocol against advanced cardiovascular disease involves the placement of an intravascular stent, an expandable wire mesh structure, into the diseased artery. The primary limitation of stenting procedures is restenosis, a complex process by which vessel blockage redevelops over a period of a few months. The deployment of a stent within an arterial segment locally damages the endothelium, the monolayer of cells lining the inner surfaces of blood vessels. As a result of this damage, both thrombotic and wound healing biological pathways are rapidly activated.

It appears that the incidence of restenosis is reduced if re-endothelialization of the damaged vascular area occurs sufficiently rapidly. Re-endothelialization requires endothelial cell proliferation and migration, and these processes are likely affected by the local hemodynamic environment. We have been using computational fluid dynamics (CFD) to study the detailed flow field in the vicinity of implanted stents. Of particular interest is the extent of flow disturbance induced by the placement of a stent within an artery. The simulations involve solving the full three-dimensional time-dependent equations governing fluid mass and momentum conservation (Navier–Stokes equations) in model straight and curved arterial segments. The effects of various parameters including stent wire thickness, inter-wire spacing, extent of vessel curvature, flow Reynolds number, and flow inlet inlet velocity profile on the occurrence of flow disturbance have been investigated.

Our results have revealed a region of flow separation and recirculation immediately downstream of model stents. The size of this recirculation flow zone increases with both the flow Reynolds number and the stent wire thickness. With flow pulsatility, this region of flow disturbance periodically appears and disappears. Interestingly, the nature of the flow disturbance downstream of a stent depends on a complex fashion on the vessel geometry, flow conditions, and stent design.

Given the wide spectrum of arterial configurations and flow conditions present in vivo, an interesting notion would be to use CFD techniques to optimize the design of stents for particular vascular sites with the goal of minimizing stent-induced flow disturbance.

August 9, 2001

New Monte Carlo Methods for Problems in Materials and Biology

Michael Mascagni
Florida State University

Email: mascagni@cs.fsu.edu

Abstract:

Probabilistic potential theory enables us to solve a large class of parabolic and elliptic partial differential equations using diffusion techniques. Here, we present new first- and last-passage Monte Carlo algorithms and show their utility in materials science and biology problems. These techniques exploit the fact that the first-passage probability function is the Green's function for the Dirchlet problem of the Laplace equation. First-passage algorithms allow the rapid simulation of diffusion using analytic or simulation-based Green's functions in rather less complicated basic geometries. This permits consideration of more complicated real geometries made up as combinations of the simple geometries where Green's functions are available. This new method is the extension of the well-known "Walk on Spheres" method.

Harnessing these first-passage algorithms, we have developed the fastest algorithms known to compute (1) the fluid permeability in overlapping, nonoverlapping, and polydispersed spherical models of random porous media, (2) the Solc-Stockmayer model with zero potential, a model of ligand binding, (3) the mean trapping rate of a diffusing particle in a domain of nonoverlapping spherical traps, and (4) the effective conductivity for perfectly insulating, nonoverlapping spherical inclusions in a matrix of conductivity.

In certain problems, such as that of computing the electrostatic charge distribution on a conductor, using the last-passage distribution is useful. Using these analogous last-passage algorithms, we have solved the test problem of computing the charge distribution on a circular two-dimensional disk in three dimensions.

Our plans for the future involve adding more surface Green's functions to our present set of known Green's functions, and the application of these techniques to more realistic problems in materials and biology.

This is joint work with Dr. Chi-Ok Hwang of Florida State University, and Dr. James Given of Angle, Inc.

Speaker's web page: http://www.cs.fsu.edu/~mascagni/

Institution web page: http://www.fsu.edu/

August 8, 2001

Effective Query Processing Over Large Multidimensional Data Sets

Sharad Mehrotra
University of California, Irvine

Email: sharad@ics.uci.edu

Abstract:

In this talk, I will present two of our recent research results in query processing over multidimensional data sets. The first result deals with answering aggregate queries in an approximate manner over very large data collections. Our approach uses selective traversal of a multiresolution aggregate (MRA) tree structure to provide 100% interval of confidence on the value of the aggregate. The algorithm works iteratively coming up with improving quality answers until some error metric is satisfied or time constraint is reached. This work is part of the Quality Aware Sensor Database Architecture (QUASAR) project, in which we explore techniques to trade quality of answer for performance in database query processing when in highly dynamic data intensive environments.

The second part of the talk will discuss local dimensionality reduction (LDR) approach that effectively overcomes the dimensionality curse problem in range and K-NN searches over highly multidimensional data sets. Unlike previous approaches, LDR exploits the local correlation in data to identify clusters that have a natural lower-dimensional representation. These clusters can then be individually indexed and searched using multidimensional data structures. This work is part of the Multimedia Analysis and Retrieval Systems (MARS) project in which we are developing a data management system that provides native and scalable support for similarity retrieval and query refinement over multimedia information.

Speaker's web page: http://www-db.ics.uci.edu/pages/people/sharad.shtml

Research web page: http://www-db.ics.uci.edu/pages/research/mars/

Institution web page: http://www.uci.edu/

July 27, 2001

Performance Analysis of MPI/OpenMP Applications Using Paraver

Jesus Labarta

European Center for Parallelism of Barcelona

Email: jesus@cepba.upc.es

Abstract:

Paraver is a visualizer for the performance of parallel programs that supports detailed analyses down to very fine granularities. Two major strengths of Paraver are:

- Flexibility and support for analyses not even planned.
- Accurate quantitative analysis support.

The OMPItrace instrumentation package for MPI, OpenMP, and mixed mode programs has been ported to IBM SP platforms in cooperation with the ACTC. Besides the process and communication activity, a wide range of performance indices derived from hardware counters can be displayed and measured. A translator from the UTE trace format to Paraver is also available.

Dimemas is a simulator useful in performance estimation of message-passing programs. Dimemas rebuilds the timing behavior of MPI programs based on simple target architecture parameters. These parameters describe basic aspects that influence application performance such as network latency and bandwidth, injection mechanism, processor performance, and predicted effect of improvement in one code section.

The presentation will give an update of the type of analyses that can be performed with Paraver and Dimemas as well as their interoperability capabilities. Examples will be mostly based on the Sweep3d application. For more information see: http://www.cepba.upc.es/paraver/ and http://www.cepba.upc.es/dimemas/.

Institution web page: http://www.cepba.upc.es/

July 26, 2001

Performance Optimization of Component-based Data Intensive Applications

Alan Sussman
University of Maryland

Email: als@cs.umd.edu

Abstract:

In the DataCutter project we have proposed the use of a component-based processing model, called filter-stream programming, for data intensive applications. In this talk I will give an overview of filter-stream programming in DataCutter, and also discuss issues related to retrieval of data stored in a distributed, high-performance, heterogeneous computing environment. The main part of the talk will discuss several of the optimization techniques we are designing and evaluating on the heterogeneous set of computational and storage Linux clusters at the University of Maryland. These optimizations include the reuse of existing groups of filters to perform a requested computation, and the transparent replication of particular filters that cause bottlenecks in a filter pipeline. I will present results from the evaluation showing that performance varies greatly depending on the assignment of filters to hosts and the number of transparent copies instantiated for each filter, and that performance is also highly dependent on the distribution of source data to host disks. In ongoing work, we are using these results to drive the construction of cost models to enable automating the placement of filters on the set of available hosts, and selecting the number of copies of each filter on each host.

Speaker's web page: http://www.cs.umd.edu/~als/

Institution web page: http://www.umd.edu/

July 25, 2001

A Multilevel Nonlinear Method

Irad Yavneh

Technion-Israel Institute of Technology

Email: irad@cs.technion.ac.il

Abstract:

Robust multigrid algorithms for linear boundary-value problems are well researched. For nonlinear problems, two general approaches are used widely: Global Linearization (GL) and Local Linearization (LL). In GL methods, the discretized problem is linearized, usually by Newton's method, and the resulting linear system is solved approximately by a linear multigrid algorithm. This is repeated iteratively. In the LL approach, the nonlinear fine-grid operator is approximated nonlinearly on the coarser grids, and explicit linearization is only performed locally, in the relaxation process. The best known of the LL methods is the so-called Full Approximation Scheme (FAS). For simple problems, the two approaches often perform similarly, but a distinct behavior is exhibited in more complicated settings, with the GL approach performing better in some cases and the LL approach in others.

We propose a Multilevel Nonlinear Method (MNM) which, we believe, will generally be at least as robust as either one of the above, and often more robust than both. The work described is at an early stage, and only a highly simplified analysis and preliminary numerical results are yet available. However, these indicate that MNM may continue to perform even as the FAS and Newton methods melt in your hand.

Institution web page: http://www.technion.ac.il/

July 24, 2001

Three-Dimensional Control-Volume Mixed Finite Element Methods and Efficient Algebraic Solvers for Distorted Hexahedral Grids

Thomas Russell
University of Colorado, Denver

Email: trussell@carbon.cudenver.edu

Abstract:

The control-volume mixed finite element (CVMFE) method is designed to calculate accurate pressure and velocity distributions in subsurface flow problems with irregular geology and heterogeneous, discontinuous, anisotropic conductivity. In two dimensions, it differs from the usual lowest-order Raviart–Thomas (RT) mixed method in its choice of velocity-vector test functions, which lead to discrete analogues of Darcy's law on control volumes. Irregular geology is modeled with logically rectangular distorted quadrilateral and hexahedral grids in 2D and 3D, respectively.

The extension from 2D to 3D raises surprising issues and dilemmas. For example, in general, the trial space obtained from the standard Piola transformation does not contain the uniform constant velocity field. The 3D formulation is presented with details about better choices of trial functions, along with a sampling of numerical results that show second-order convergence when the exact solution is not singular. Also presented, in the context of the RT method with a forthcoming extension to CVMFE, is an efficient solver for the 3D equations. This uses a convenient basis for the divergence-free velocity functions to reduce the equations to a symmetric positive-definite system of smaller size than has been possible previously. With conjugate gradients preconditioned by overlapping domain decomposition, this enables the 3D equations to be solved with effort comparable to that for other methods.

Institution web page: http://www.cudenver.edu/

July 19, 2001

Two Performance Studies in Support of the NSF/DOE Coupled Climate Model

Patrick Worley
Oak Ridge National Laboratory

Email: worleyph@ornl.gov

Abstract:

We describe two performance studies that were undertaken in support of the development and optimization of the NSF/DOE Coupled Climate Model. The first study motivates the design of a new data structure for the column physics routines used in the atmospheric model. The new data structure allows index ordering and loop ordering and length to be set at compile or runtime. We use empirical studies on Compaq AlphaServer SC, IBM SP, and NEC SX5 systems to demonstrate the performance gains that are made possible by this type of flexibility.

The second study is an evaluation of the "allreduce" collective communication operator on the Compaq AlphaServer SC and IBM SP. We compare the performance of the vendor-supplied MPI_ALLREDUCE with that of a number of different algorithms for a variety of vector lengths, processor counts, and assumptions on how the allreduce is used in an application code. Best observed performance is used to identify good parallel algorithms and to compare performance between systems.

Speaker's web page: http://www.csm.ornl.gov/~worley/

Research web page: http://www.csm.ornl.gov/evaluation/

Institution web page: http://www.ornl.gov/

July 18, 2001

A Fast Poisson Solver for Irregular 3D Regions

Anne Greenbaum
University of Washington

Email:
greenbau@math.washington.edu

Abstract:

We describe a fast method for solving Poisson's equation on irregular 3D regions. The method consists of two steps, each carried out with a fast Poisson solver (or a multigrid solver) applied on an embedding cube.

The first step is to evaluate a volume integral in order to obtain some solution to the inhomogeneous equation. Using known jumps in the second derivatives of the volume integral, one can determine the right-hand side for a discrete version of Poisson's equation on a uniform grid throughout an embedding cube. This equation can then be solved with an FFT or a multigrid method. Next, one solves a boundary integral equation to determine an integral representation of the solution to Laplace's equation with appropriate boundary conditions. Knowing this integral representation, one can compute the discontinuities in the solution and its derivatives across the boundary of the region and use these to determine the right-hand side for another fast Poisson solve on the embedding cube. The total cost is that of two fast Poisson solves on a cube [O(N log N) or O(N), depending on whether an FFT or a multigrid method is used] plus the cost of solving a boundary integral equation [O(m^2) or O(m), where m is the number of boundary discretization points, depending on whether the dense *m* by *m* matrix is assembled or the fast multipole method is used to compute matrix-vector products].

A software package implementing these ideas is currently being constructed. The algorithms are amenable to parallelization, and specific problems have been solved using multiple processors on an IBM SP-2.

Speaker's web page: http://www.washington.edu/~greenbau/

Institution web page: http://www.washington.edu/

July 18, 2001

Parallel Data Access in Heterogeneous Environments with Storage Tank

Randal Burns
IBM Almaden Research Center

Email: randal@almaden.ibm.com

Abstract:

We describe how the Storage Tank distributed file system manages platform and network heterogeneity in a large-scale computing environment. Storage Tank provides shared access from multiple operating systems to fibre channel, iSCSI, or object-based network attached storage. This allows multiple computers or clusters to share data without degrading performance at compute nodes.

We also propose data placement algorithms that de-cluster I/O and place data near where they are used in heterogeneous networks. Declustering will help Storage Tank achieve aggregate write rates much greater than the capacity of any single switch or link in the system. Storage Tank declusters data on a per-file basis using a policy engine, and within a single file through distributed storage pools, both of which adapt to changing workloads and component failures.

Speaker's web page: http://www.cse.ucsc.edu/~randal/

Institution web page: http://www.almaden.ibm.com/

July 17, 2001

TAU Performance System: Developments and Evolution

Allen Malony
University of Oregon

Email: malony@cs.uoregon.edu

Abstract:

The TAU performance system offers a robust and portable performance instrumentation, measurement, and analysis framework and toolkit for its implementation. TAU is based on a general parallel computation model and is targeted to large-scale, multi-thread shared memory and distributed memory message-passing parallelism. With over eight years of research and development, the TAU performance system is mature and is now being applied in several ASCI environments.

This talk gives some highlights of current TAU developments and application. New tools include a new Java-based parallel profile display tool and a new Fortran90 front-end for our Program Database Toolkit. We will present case studies of TAU's use for multi-threaded and mixed-mode performance measurement, and also our work with the Utah Uintah parallelism framework being applied in the C-SAFE ASCI effort. This latter effort includes an initial version of an experimental management system that will be described. Recently, we have proposed and prototyped an OpenMP performance API. Its design and implementation will be presented in the talk.

The TAU performance system will evolve in several important directions that are a natural extension of current capabilities as well as integration with other key technologies. We will describe these future activities.

Speaker's web page: http://www.cs.uoregon.edu/~malony/

Research web page: http://www.cs.uoregon.edu/research/paracomp/tau/

Institution web page: http://www.uoregon.edu/

July 16, 2001

Parallel and Robust Multigrid Techniques on Structured Grids

Ignacio Llorente
Complutense University, Madrid

Email: llorente@eucmax.sim.ucm.es

Abstract:

Research presented at this seminar, carried out by DACYA/UCM (Departamento de Arquitectura de Computadores y Automatica/Universidad Complutense de Madrid) in collaboration with ICASE (Institute of Computer Applications in Science and Engineering), focuses on the development of parallel and robust multigrid techniques on structured grids. We present results from the parallel implementation of different robust alternatives to solve the 3D diffusion-convection operator in highly non-aligned anisotropic problems. This study was used as a starting point for the development of a robust solver for the 3D Navier–Stokes equations. The results include both algorithmic and architectural properties (convergence rate, efficiency, scalability, and cache-oriented issues) on two different architectures, namely a low-cost distributed memory architecture (CORAL Beowulf-class system) and a cc-NUMA shared memory machine (SGI Origin 2000).

Speaker's web page: http://www.dacya.ucm.es/nacho/

Research web page: http://www.dacya.ucm.es/nacho/research/research.html

July 13, 2001

Hybrid Intelligent Systems for Industrial Data Analysis

Arthur Kordon
The Dow Chemical Company

Email: akordon@dow.com

Abstract:

Industry has recently significantly increased its efforts toward effective data analysis. Many companies even look at industrial data analysis as one of the main driving forces for establishing strategic competitive advantage. Hybrid intelligent systems, which integrate various techniques like neural networks, genetic programming, support vector machines, fuzzy logic, knowledge acquisition, and so forth, can play a significant role in this process. The objective of the presentation is to summarize the research topics and implementation issues of hybrid intelligent systems for industrial data analysis.

The first part of the presentation is focused on the special features and issues of industrial data analysis. In contrast to research-driven data analysis where the objective is to "transfer data into knowledge," the goal of industrial-driven analysis is to "transfer data into value." The difference in the objective function has an inevitable effect on the whole strategy for data analysis. Of special importance is the constraint on the level of complexity of delivered models (i.e., there is no need to increase complexity if the economic objective is accomplished). This strategy pushes industrial data analysis toward simple and robust empirical solutions that can be achieved with minimal time and training efforts.

The following special features of industrial data analysis are discussed briefly: hidden effect of operators' intervention, "curse of closed loops" effect, multiple time scale data analysis, and real-time pressure. The current state of the art and some issues of industrial data analysis, such as non-adaptive nature of the delivered empirical models and lack of flexible integration of the most appropriate data analysis methods and knowledge sources, are covered.

The second part of the presentation is devoted to the use of hybrid intelligent systems in analyzing industrial data. The "classic" hybrid intelligent systems of neural nets, fuzzy logic, genetic algorithms, and rule-based knowledge are compared with the "modern" integration of stacked analytic neural nets (not based on back-propagation), support vector machines, and genetic programming. The open research topics for development of modern hybrid intelligent systems (e.g., kernel selection for support vector machines, VC-dimension for analytic neural nets, integrating knowledge in genetic programming, etc.), as well as the main application areas like adaptive soft sensors, real-time design of experiments, and data-driven fundamental model emulators are discussed. The potential of hybrid intelligent systems for industrial data analysis is shown with several applications in The Dow Chemical Company (robust soft sensor in a chemical reactor, automating operating discipline in a large-scale chemical plant, using genetic programming for minimizing the design of experiments, and data-driven emulators for complex fundamental models).

The final part of the presentation defines the main industrial needs for improving data analysis. On the top of the list is the need to integrate technical and business-related data analysis. This is of special importance to global companies operating in the real-time financial environment of the global market. Another key factor for successful mass scale implementation of this approach is the development of autonomous process data analysis in real time.

Institution web page: http://www.dow.com/

July 12, 2001

Implementing Parallel Shortest Path for Parallel Transportation Applications

Michelle Hribar
Pacific University

Email: hribarm@pacificu.edu

Abstract:

Shortest path algorithms are used by many different applications; network routing, VLSI design, and transportation algorithms are just a few. Shortest path computation can be very expensive. Algorithms can take up to *O(n**2)*, where *n* is the size of the network (number of nodes or vertices). As the size of the network grows, this shortest path computation can dominate the execution time of the application. Therefore, parallel processing can provide the compute power required to quickly solve large problems.

Parallel shortest path algorithms are a well-studied topic, particularly in the area of theoretical development. In this talk, I will discuss an experimental study of parallel shortest path algorithms that focuses on understanding the performance impact of different implementation issues. In particular, I will focus on three implementation areas: (1) choice of shortest path algorithm, (2) termination detection, and (3) network decomposition. I will show how implementation decisions in all three areas have a large impact on the communication and convergence of parallel shortest path algorithms. In general, we find that communicating the most information at a time results in the best convergence and overall execution time. This is contrary to most scientific applications where it is optimal to minimize communication. I will present these results in the context of a transportation application that solves for traffic equilibrium.

Speaker's web page: http://zeus.cs.pacificu.edu/~hribarm/

Institution web page: http://www.pacificu.edu/

July 9, 2001

Synthesizing Sounds from Physically Based Motion

James O'Brien
University of California, Berkeley

Email: job@cs.berkeley.edu

Abstract:

This talk describes a technique for approximating sounds that are generated by the motions of solid objects. The technique builds on previous work in the field of physically based animation that uses deformable models to simulate the behavior of the solid objects. As the motions of the objects are computed, their surfaces are analyzed to determine how the motion will induce acoustic pressure waves in the surrounding medium. The technique computes the propagation of those waves to the listener and then uses the results to generate sounds corresponding to the behavior of the simulated objects.

Speaker's web page: http://www.cs.berkeley.edu/~job/

Institution web page: http://www.berkeley.edu/

July 6, 2001

A Scalable Hierarchical Algorithm for Unsupervised Clustering

Daniel Boley
University of Minnesota

Email: boley@cs.umn.edu

Abstract:

Top-down hierarchical clustering can be done in a scalable way. Here we describe a scalable, unsupervised clustering algorithm designed for large datasets from a variety of applications. The method constructs a tree of nested clusters top-down, where each cluster in the tree is split according to the leading principal direction. We use a fast principal direction solver to achieve a fast overall method. The algorithm can be applied to any dataset whose entries can be embedded in a high-dimensional Euclidean space, and takes full advantage of any sparsity present in the data. We show the performance of the method on text document data, in terms of both scalability and quality of clusters. We demonstrate the versatility of the method in different domains by showing results from text documents, human cancer gene expression data, and astrophysical data. For that last domain, we use an out-of-core variant of the underlying method, which is capable of efficiently clustering large datasets using only a relatively small memory partition.

Speaker's web page: http://www.cs.umn.edu/~boley/

Institution web page: http://www.umn.edu/

July 6, 2001

Dynamic Meshes, Dynamic Interfaces, and Hemodynamics

Omar Ghattas
Carnegie Mellon University

Email: oghattas@cs.cmu.edu

Abstract:

Many important phenomena in science and engineering, including our motivating problem of microstructural blood flow, can be modeled as flows with dynamic interfaces. The major challenge faced in simulating such flows is resolving the interfacial motion. Lagrangian methods are ideally suited for such problems, since interfaces are naturally represented and propagated. However, the material description of motion results in dynamic meshes, which become hopelessly distorted unless they are regularly regenerated. Lagrangian methods are particularly challenging on parallel computers, because scalable dynamic mesh methods remain elusive.

I will present a parallel dynamic mesh Lagrangian method for flows with dynamic interfaces that we have been developing at CMU. We take an aggressive approach to dynamic meshing by triangulating the propagating grid points at every time step using a scalable parallel Delaunay algorithm. Contrary to conventional wisdom, I will provide evidence that the costs of the geometric components (triangulation, coarsening, refinement, and partitioning) can be made small, relative to the flow solver. For example, in a 2D simulation of 10 interacting viscous cells with 500,000 unknowns on 64 processors of a Cray T3E, dynamic meshing consumes less than 5% of a time step. Moreover, our experiments on up to 128 processors show that the computational geometry scales about as well as the flow solver.

I will discuss the application of our dynamic mesh Lagrangian method to microstructural simulation of blood flow, which is essentially a problem in modeling the interaction of fluid-solid mixtures. The model is termed "microstructural" because it distinguishes the fluid (blood plasma and hemoglobin) from the solid (cell membrane) at micron scales and computes the momentum exchange between them. This is in contradistinction to typical macroscopic models that treat blood as a homogeneous viscous medium with phenomenological incorporation of cellular effects. I will conclude with a discussion of the prospects for microstructural modeling of blood flow at scales of interest in the design of artificial heart devices.

This work is jointly conducted with graduate student Ivan Malcevic, CMU colleagues Guy Blelloch, Gary Miller, and Noel Walkington, and University of Pittsburgh Medical Center collaborator Jim Antaki.

Speaker's web page: http://www.cs.cmu.edu/~oghattas/omar.html

Institution web page: http://www.cmu.edu/

June 25, 2001

Relaxation Schemes and Approximate Riemann Solvers for Nonlinear Hyperbolic Problems with Variable Coefficients

Randall LeVeque
University of Washington

Email: rjl@amath.washington.edu

Abstract:

A nonlinear system of *m* conservation laws can be rewritten as a linear system of *2m* equations, with a stiff source term that drives the solution towards that of the original equation. Relaxation schemes, numerical methods based on this idea, were introduced by Jin and Xin and have been used for many problems. These methods can be reinterpreted as a way to define an approximate Riemann solver for the original problem. This viewpoint suggests a broader class of "relaxation Riemann solvers" that approximate the solution to a system of *m* equations using *2m* waves. Nonlinear conservation laws with spatially varying flux functions can often be solved using an approximate Riemann solver of this form. One interesting application is to nonlinear elastodynamics in a laminate of thin layers, and hence rapidly varying material parameters. An initially smooth disturbance starts to sharpen into a shock but becomes oscillatory and appears to break up into a train of solitons due to the dispersive nature of the heterogeneous material.

Speaker's web page: http://www.amath.washington.edu/~rjl/

Research web page: http://www.amath.washington.edu/~claw/

Institution web page: http://www.washington.edu/

June 19, 2001

Efficient Adaptive Simplification of Massive Meshes for Visualization

Eric Shaffer

University of Illinois, Urbana-Champaign

Email: shaffer1@cs.uiuc.edu

Abstract:

The growing availability of massive polygonal models, and the inability of most existing visualization tools to work with such data, has created a pressing need for memory efficient methods capable of simplifying very large meshes. In this talk, we will present a method for performing adaptive simplification of polygonal meshes that are too large to fit in-core.

Our algorithm performs two passes over an input mesh. In the first pass, the model is quantized using a uniform grid, and surface information is accumulated. This sampling is then used to construct a BSP-Tree in which the partitioning planes are determined by the surface information. In the final pass, the original vertices are clustered using the BSP-Tree, yielding an adaptive approximation of the original mesh. The BSP-Tree describes a natural simplification hierarchy, making it possible to generate a progressive transmission and construct level-of-detail representations. In this way, the algorithm provides some of the features associated with more expensive edge contraction methods while maintaining greater computational efficiency. In addition to performing adaptive simplification, our algorithm exhibits output-sensitive memory requirements and allows fine control over the size of the simplified mesh.

Speaker's web page: http://bugle.cs.uiuc.edu/People/shaffer1/

Institution web page: http://www.uiuc.edu/

June 15, 2001

Multiscale Computational Tools for Elementary Particles

Achi Brandt
Weizmann Institute of Science

Email: achi.brandt@weizmann.ac.il

Abstract:

Current calculations of quantum fields are extremely expensive and obtain only very poor predictions of elementary particle properties. The enormous cost is the product of several separate complexity factors, which will be explained. In a series of recent studies of simple models, it has been shown that each of these complexity factors can be eliminated or drastically reduced by several new types of multiscale algorithms, including those described below.

A new algebraic multigrid (AMG) technique can solve Dirac Equations on highly disordered gauge fields at linear complexity. The "bootstrap AMG," an approach for fast-solving general discretized PDE systems, will be described in detail.

A special multiscale structure for describing the main part of the inverse of a differential operator allows very fast self-updating upon local changes in the operator. This permits instant updating of the inverse and the operator dominant during Monte Carlo simulations of the gauge field, yielding fast calculations of the fermionic interactions.

A new Monte Carlo approach that combines renormalization and multigrid ideas has been shown to eliminate both the critical slowing down and volume factors, implying that thermodynamic limits can generally be calculated to accuracy *e* in O(e**(-2)) computer operations. Related renormalization-multigrid (RMG) procedures are being developed for fast simulations in statistical mechanics and molecular dynamics.

Speaker's web page: http://www.wisdom.weizmann.ac.il/~achi/

Institution web page: http://www.weizmann.ac.il/

June 12, 2001

ACTC Tools for Application Performance Analysis of Scientific Programs

Luiz De Rose
IBM TJ Watson Research Center

Email: laderose@us.ibm.com

Abstract:

In this talk I will present an overview of the application performance analysis tools under development at the Advanced Computing Technology Center for tuning and optimization of applications running on the IBM platforms. These tools were designed to help users understand the behavior of applications on complex parallel environments. They are based on dynamic instrumentation, access to hardware performance counters, hints to allow users to correlate the behavior of the application to hardware components, and careful mapping of performance data to source code constructs. The presentation will focus on the Hardware Performance Monitor (HPM) Toolkit and the Simulator Guided Memory Analyzer (SiGMA).

The HPM Toolkit consists of libraries and utilities to support performance data capture, analysis, and presentation of hardware performance metrics from applications written in Fortran, C, and C++, executing on sequential or parallel systems. These metrics allow users to correlate the behavior of the application to one or more of the components of the hardware, providing hints that help users to identify causes of performance problems.

SiGMA is a data-centric tool under development that will be able to predict performance and identify bottlenecks, problems, and inefficiencies in a program, due to data layout in the memory hierarchy, as well as propose solutions to improve program performance in current and future architectures.

Institution web page: www.research.ibm.com/actc/

June 8, 2001

The Aqua Approximate Query Answering System

Phillip Gibbons
Bell Laboratories

Email:
gibbons@research.bell-labs.com

Abstract:

Exploratory data analysis (also called OLAP, decision support, and data mining) can be quite time consuming for large data recording and warehousing environments (e.g., sales transaction databases, call detail repositories, customer service historical data). Responses to queries can take minutes to hours, even on top-of-the-line database systems, due to the large amount of computation and disk I/O required to compute an exact answer over a large database.

Aqua is the first system to provide fast, highly accurate approximate answers for a broad class of queries arising in data warehousing environments. Aqua answers queries using precomputed synopses of the data. Approximate answers with error guarantees (e.g., 384K +/- 1K) are provided by rewriting the user query over the synopses and executing the new query. Aqua provides answers in orders of magnitude less time than it takes to compute an exact answer. To avoid costly overheads, Aqua synopses are incrementally maintained in the presence of database updates. Besides the overall system architecture, the key to Aqua is the design and analysis of its synopses. We give examples of novel synopses developed for Aqua; these are also applicable to other "data streams" contexts such as network monitoring.

Aqua is a middleware software solution supporting standard interfaces based on ODBC, to work with any report generating software in the front-end and any commercial DBMS in the back-end. It is currently available for the Windows and Solaris (UNIX) platforms.

Speaker's web page: http://www1.bell-labs.com/user/pbgibbons/

Research web page: http://www.bell-labs.com/project/aqua/

Institution web page: http://www.bell-labs.com/

June 6, 2001

Memory Profiling on Shared-Memory Multiprocessors

Jeff Gibson
Stanford University

Email: jeffg@powderkeg.stanford.edu

Abstract:

Tuning application memory performance can be difficult on any system, but it is particularly so on shared-memory multiprocessors due to the implicit nature of communication and the unforeseen interactions among the processors. Tools, called memory profilers, that allow the user to map memory behavior back to application data structures can be invaluable aids to the programmer. Unfortunately, memory profiling is difficult to implement efficiently since most systems lack the requisite hardware support.

This work describes several techniques that can be used for efficient memory profiling. Each requires differing levels of hardware support, either on the compute processor or on the system node controller. We prototype each system on the versatile FLASH multiprocessor and present measurements of both overhead and accuracy for each method.

Speaker's web page: http://www-flash.stanford.edu/~jeffg/

Research web page: http://www-flash.stanford.edu/

Institution web page: http://www.stanford.edu/

June 4, 2001

Efficient Shared Memory Support in NUMA-based Cluster Environments

Martin Schulz
Technical University, Munich

Email: schulzm@in.tum.de

Abstract:

NUMA-based clusters provide an efficient and cost-effective platform for parallel computing. However, a comprehensive software infrastructure with support for all major parallel programming paradigms is required in order to fully exploit their potentials. Such an infrastructure has been developed within the SMiLE project (http://smile.in.tum.de/) for commodity clusters based on the NUMA interconnection fabric SCI (Scalable Coherent Interface).

While support for message passing can be achieved on top of these architectures in a straightforward manner, the implementation of a Shared Memory programming framework imposes additional challenges. Most severely are the semantic gap between the global physical memory provided by the NUMA hardware and the global virtual memory required for Shared Memory programming and the need for an explicit cache and consistency management.

The HAMSTER (Hybridism-based Adaptive and Modular Shared memory archeTEctuRe, http://hamster.in.tum.de/) framework aims to overcome these challenges and thereby contribute to a multi-paradigm software infrastructure efficiently supporting both message passing and shared memory. Its core component is a Hybrid-DSM system combining the NUMA hardware capabilities with a lean software component for memory and consistency control. A prototype has been implemented within SMiLE and has been evaluated using several applications, including two real-world codes from the area of nuclear medical imaging. The results of these experiments are very encouraging and prove both the feasibility and the efficiency of the proposed approach.

Speaker's web page: http://wwwbode.in.tum.de/~schulzm/

Institution web page: http://wwwbode.in.tum.de/

June 4, 2001

ARIMA Time Series Dynamic Modeling and Forecasting for Adaptive I/O Prefetching

Nancy Tran

University of Illinois, Urbana-Champaign

Email: n-tran1@cs.uiuc.edu

Abstract:

Variability in application data access patterns, together with fluctuations in processor loads and network delays, make input/output (I/O) request arrivals irregular and bursty. Such bursty patterns, combined with transfer-limited storage devices, create a major I/O bottleneck. One approach at relieving this bottleneck is to anticipate future I/O bursts, adaptively prefetching data before it is needed. To be effective, prefetches must be initiated early enough to avoid I/O stalls, but not too early, to minimize unnecessary cache buffer replacement. In this talk, we will introduce the use of time series analysis to dynamically model and forecast diverse I/O temporal patterns during application execution. The forecasts are then used to prefetch read requests during computation intervals to hide I/O latency.

We will describe the design of a "just-in-time" prefetcher. This prefetcher combines time series forecasts with the Pablo group's Markov model spatial forecasts to adaptively determine when, what, and how much to prefetch. Experimental results with I/O intensive scientific codes show improvement in total execution time and cache misses over the traditional Unix file system.

Research web page: http://www-pablo.cs.uiuc.edu/

Institution web page: http://www.uiuc.edu/

May 31, 2001

Immediate Access to Highly Complex, Remote, 4D Models

Jack Snoeyink and Jarek Rossignac

University of North Carolina, Chapel Hill, and Georgia Institute of Technology

Abstract:

Part I: Compression and progressive transmisison of meshes of triangles and tetrahedra (J. Rossignac)

After developing IBM's "Topological Surgery" technique with G. Taubin, which is now the core of the MPEG-4 standard for 3D compression, Jarek Rossignac has focused his research efforts on a simpler compression scheme, called "Edgebreaker." Edgebreaker holds the record for the lowest guaranteed upper bound for encoding planar triangle graphs (1.8 bits per triangle). In practice, it compresses the connectivity of triangle meshes down to about a bit per triangle. A new implementation of Edgebreaker, developed in collaboration with A. Szymczak, A. Safonova, H. Lopes, and V. Coors, works as a simple state machine. The detailed code fits on three pages and uses, as sole data structure, two arrays of integers, called the Corner Table.

The current implementation, which is publicly available, uses a parallelogram prediction but can be easily combined with other schemes for compressing vertex location. Both the Topological Surgery and Edgebreaker have received Best Paper awards. Rossignac and his colleagues have also developed a similar approach, called Grow&Fold, for compressing tegrahedral meshes, such as those used for finite element analysis. Rossignac has also developed progressive techniques for transmitting compressed versions of both triangle meshes and of tegrahedral meshes (with R. Pajarola and A. Szymczak), such that the client may download low resolution models first and then, when necessary, refine them to higher accuracy by downloading compressed upgrades.

Part II: Isosurfaces and contour trees (J. Snoeyink)

Scientific simulations produce data in the form of sample points with intensity values, which are often visualized by drawing level sets or isosurfaces (surfaces of points with the same intensity). One tool that can help in choosing threshold values for interactive exploration of such data is the contour tree, which has been developed by a number of people including van Kreveld, Bajaj, and Pascucci. Contour trees encode the evolution of isosurfaces as the threshold parameter varies. It can be used to compute seed sets for tracing whole or partial isosurfaces, to determine important values for the threshold parameter, and to support flexible contouring. Carr, Snoeyink, and Axen developed a simple algorithm to compute a contour tree without computing the set of all isosurfaces, which enables the use of the contour tree as a data analysis tool.

Part III: Towards a real-time remote visualization of time-dependent 3D simulations (J. Rossignac)

Engineering and scientific simulations conducted at high-performance computing centers produce hundreds of gigabytes of data, which represent the evolution of several variables over a 4D space-time domain. New compression and progressive transmission techniques are needed to support the interactive

(continued next page)

Abstract (continued):

Jack Snoeyink and Jarek Rossignac

University of North Carolina, Chapel Hill, and Georgia Institute of Technology

exploitation of this data. In collaboration with J. Snoeyink and his colleagues at UNC and with A. Szymczak and S. Menon (Aerospace) at Georgia Tech, Jarek Rossignac has been exploring new approaches for the compression, progressive transmission, and interactive visualization of such data sets. The techniques under consideration view the data as a hyper-terrains in five dimensions.

The natural sampling of the hyper-terrain on a regular 4D space-time grid may be concisely encoded at several levels of resolution using higher-dimensional wavelets or other predictive schemes coupled with variable-length coding. However, in interactive applications where selected subsets of the data must be accessed at specific resolutions, regular representations may prove less effective than pentatope meshes produced through an adaptive simplification process or than semi-regular multi-resolution representations. We plan to evaluate all three approaches in the context of a client/server system, where the operator uses two parameters (say pressure P and time T) to interactively control a color-coded iso-surface S(T,P) of all 3D points that take pressure P at time T.

The relevance of each iso-surface, measured for example in terms of its surface area, curvature integral, or number of connected components, may be displayed as the height or color of a P-T control terrain, upon which the (P,T) locations of previously inspected iso-surfaces may be traced and annotated. The global view of the entire dataset provided by the control terrain may help focus interactive inspection on the relevant subsets of the data. We plan to extend to these 4D models the compression and progressive transmission techniques that we have developed for 3D meshes, so that a crude approximation of an initial isosurface S(P,T) may be quickly downloaded and then refined in real time either to follow user-controlled changes in P or T, or to increase its accuracy during navigation pauses.

Part IV: An interface prototype (J. Snoeyink)

We have implemented an initial prototype for the user interaction on the client with Lutz Kettner, a postdoctoral researcher at UNC Chapel Hill. We illustrate this prototype on data sets from combustion simulation, fluid flow, pollution studies, and heat convection in the earth. The data dimensions on the client side are partitioned into a 3D viewer for isosurfaces and a 2D control plane where each point selects a particular isosurface in 3D. Annotations in the control plane help the user to navigate the volume data. We use the number of connected components of the isosurface as an example of an automatically generated annotation that we can compute efficiently using contour trees. We also provide a small pre-computed preview window that is shown continuously at the current cursor position in the control plane. The small preview facilitates a fast overview; selecting the current location in the control plane creates the corresponding detailed view in the 3D viewer for isosurfaces.

May 22, 2001

Higher-Order Semi-Implicit Projection Methods for the Incompressible Navier–Stokes Equations

Michael Minion
University of North Carolina

Email: minion@amath.unc.edu

Abstract:

This talk presents a detailed discussion of the appropriate boundary conditions for semi-implicit projection methods for the incompressible Navier–Stokes equations. In particular, the boundary conditions for the intermediate velocity (the solution of the implicit treatment of the momentum equation) will be analyzed. The relationship between the form of the momentum equation, the boundary conditions for the intermediate velocity, and the pressure equation for second-order projection methods will be made clear. Also, a discussion of higher-order methods based on Spectral Deferred Corrections will be presented.

Institution web page: http://www.northcarolina.edu/

May 18, 2001

Parallel Communication & MPI-I/O Benchmarking and Automatic MPI Profiling

Rolf Rabenseifner
University of Stuttgart

Email: rabenseifner@hlrs.de

Abstract:

The talk will start with an overview on German HPC centers and on the High Performance Computing Center Stuttgart (HLRS) that has systems totaling more than 700 GFlops (Rmax), including new vector machines from NEC and Hitachi.

The second part of the talk describes our automatic counter profiling of all MPI applications on our Cray T3E-900/512. A detailed summary of the hardware performance counters and the MPI calls of any MPI production program is gathered during execution and written on a special syslog file. User specific summaries are sent by e-mail to each user each week. Summary results of profiling our users for six months are presented.

The third part of the talk describes two MPI benchmarks created to characterize the balanced system performance of high-performance clusters and supercomputers: b_eff, the communication specific benchmark examines the parallel message passing performance of a system, and b_eff_io, which characterizes the effective I/O bandwidth. Both benchmarks have two goals: (a) to get a detailed insight into the performance of different parallel communication and I/O patterns, and (b) to obtain a single average bandwidth number. Results are given for IBM SPs, Cray T3E, NEC SX-5, and Hitachi SR 8000.

Speaker's web page: http://www.hlrs.de/people/rabenseifner/

May 16, 2001

Cartesian Grid Embedded Boundary Finite Difference Methods for Partial Differential Equations in Irregular Geometries

Philip Colella

Lawrence Berkeley National Laboratory

Email: colella@colella2.lbl.gov

Abstract:

In this talk, we give an overview of a set of methods being developed for solving classical PDEs in irregular geometries, or in the presence of free boundaries. In this approach, the irregular geometry is represented on a rectangular grid by specifying the intersection of each grid cell with the region on one or the other side of the boundary. This leads to a natural conservative discretization of the partial differential operator on either side of the boundary, with the solution itself being defined on a rectangular mesh. This representation of finite differences in irregular geometries is very appealing for a variety of reasons. The problem of grid generation has already been solved. More generally, the regular geometric structure of rectangular grids simplifies a variety of issues, including control of the truncation error, the development of efficient iterative methods based on multigrid, the introduction of adaptive meshes, and the coupling to other physical submodels, such as particles or radiation. In order to understand these methods, we want to place them in a more systematic numerical analysis setting. Among the tools we will use are truncation error analysis using modified equation models with singular right-hand sides, and the development of specialized discretizations that maintain uniform stability and/or conditioning in the presence of arbitrarily small control volumes.

Speaker's web page: http://seesar.lbl.gov/anag/staff/colella/colella.html

Research web page: http://seesar.lbl.gov/anag/

Institution web page: http://www.lbl.gov/

May 15, 2001

UPS: Simplifying the Performance on Clusters of SMPs

Richard Barrett
Los Alamos National Laboratory

Email: rbarrett@lanl.gov

Abstract:

Large-scale physics simulations must be designed so that they execute effectively in a distributed memory, parallel processing computing environment. Attempts have been made to hide this environment from the application writer (for example High Performance Fortran (HPF)). Unfortunately, performance has been disappointing for many complex programs. Thus far, application writers seeking high performance have been forced to use explicit message-passing software, which has sometimes been referred to as the assembly language of parallel computers. UPS is designed to provide an approach in between these two extremes.

UPS requires that the programmer confront the computing environment, while abstracting away the complexities necessarily associated with it. The result is that the code physicist can use a simplified coding style, natural to the application, which minimizes the time spent moving the data among the distributed processes. The UPS programming model focuses on applications that can be posed using collective operations. In the basic sense this includes broadcasts and reductions of distributed data and barriers. From a higher perspective, this includes gather/scatter capabilities, which is especially useful for unstructured grid applications. Linear solver capabilities (Krylov and AMG methods) have also recently been added.

In this talk I will describe the functionality provided by the UPS library, will discuss our plans for the future, specifically within the context of anticipated computing platforms, and will attempt to determine if collaboration between us would be fruitful.

Speaker's web page:

http://www-xdiv.lanl.gov/XCI/PROJECTS/UPS/rbarrett/rbarrett.html

Research web page: http://www-xdiv.lanl.gov/XCI/PROJECTS/UPS/

Institution web page: http://www.lanl.gov/

May 11, 2001

Approximate Query Answering Using Wavelets

Rajeev Rastogi
Bell Laboratories

Email:
rastogi@research.bell-labs.com

Abstract:

Approximate query processing has recently emerged as a viable, cost-effective solution for dealing with the huge data volumes and stringent response time requirements of today's Decision Support Systems (DSS). Most work in this area, however, has so far been limited in its query processing scope, typically focusing on specific forms of aggregate queries. Furthermore, conventional approaches based on sampling or histograms appear to be inherently limited when it comes to approximating the results of complex queries over high-dimensional DSS data sets.

In this talk, we propose the use of multi-dimensional wavelets as an effective tool for general-purpose approximate query processing in modern, high-dimensional applications. Our approach is based on building wavelet-coefficient synopses of the data and using these synopses to provide approximate answers to queries. We develop novel query processing algorithms that operate directly on the wavelet-coefficient synopses of relational data, allowing us to process arbitrarily complex queries entirely in the wavelet coefficient domain. This, of course, guarantees extremely fast response times since our approximate query execution engine can do the bulk of its processing over compact sets of wavelet coefficients, essentially postponing the expansion into relational tuples until the end-result of the query.

Speaker's web page: http://www.bell-labs.com/user/rastogi/

Research web page: http://www.bell-labs.com/project/serendip/

Institution web page: http://www.research.bell-labs.com/

May 9, 2001

Parallel Adaptive Multigrid Methods with Application to LES

Stefan Lang and Sandra Nagele
University of Heidelberg

Abstract:

Advanced parallel applications based on the message passing paradigm are difficult to design and implement, especially when solution adaptive techniques are used and three-dimensional problems on complex geometries are faced. Many of these difficulties are addressed inside the UG software platform, e.g., dynamic load migration and load balancing, parallel grid adaption, basic grid management, and parallel IO. The main idea during the design process of UG, a multigrid code for the computation of partial differential equations, was to find proper abstractions for each of the different functionality parts of a parallel, adaptive, and unstructured software system. This ensures a maximal degree of code reuse. Therefore, the treatment of various partial differential equations is possible without superfluous coding. In a parallel context grid, adaption involves the need to rebalance the computational load. This stage involves both determining a new load balancing and dynamically redistributing the objects of the grid parts. A key feature of UG is the capability to dynamically migrate grid objects between the processors during run time. This difficult task is supported by DDD (dynamic distributed data), a new parallel programming model. Parallel adaptive calculations of complex real-world problems, flow in porous media and elastoplasticity, based on multigrid methods are presented.

The use of multigrid methods in conjunction with a Large Eddy Simulation (LES) is very compelling since both are based on multiple scales. LES is a turbulence model that resolves large turbulent scales and models the small ones. The scale separation is performed by applying a spatial convolution operator (filter operator) to the incompressible Navier–Stokes equations. As the filter operator, a top hat filter is applied with a grid-dependent support size. Hence, the application of the filter results in a locally varying average in space. The LES model itself also depends on the grid size since the filtering process removes all subgrid scales. Some special dynamic LES models have been developed by various researchers, which apply two filters with different support size at each point of the domain to the governing equation system. Through comparison of the two different large-scale resolutions, the model term can be specified locally. This is similar to the multigrid cycle where the defect is restricted to the coarser grid and higher frequencies are removed. Another property of dynamic models is their ability to adjust themselves to local flow structures. This adaptivity is very useful since in some regions of the domain the flow can be laminar and a turbulence model is not necessary at all. Hence dynamic models were used in the simulations.

Since unstructured grids are used, the filter width varies in the solution domain. Using unstructured grids the resolution of the turbulent scales can be increased locally to decrease the modeling effort. In the neighborhood of walls, for example, it is possible to use a smaller grid size. By this grid adaptation the local flow structures can be better resolved and less modeling with less modeling error is necessary.

A Krylov subspace method with linear multigrid as preconditioner is used to solve the linearized system. Within the multigrid cycle it is important to separate modeling and solving in the sense that only on the finest grid the modeled part of the equations is determined as described above. This is necessary for an appropriate modeling of the subgrid turbulent scales. Afterwards the model term is restricted to the coarse grids and a standard linear multigrid cycle can be used. This solution strategy was applied to different flow problems, which will be presented.

Institution web page: http://www.uni-heidelberg.de/

May 8, 2001

Low Mach Number Compressible Flow Simulations

Achim Gordner
University of Heidelberg

Abstract:

Flow simulations that take into account variable density are of importance in modeling laminar flow combustion. In recent laminar flame simulations, we used an approach based on rigorous low Mach number asymptotics. In the zero Mach number limit the acoustic modes are eliminated. However, for some applications it is necessary to take the effects of acoustic waves into account. Since low Mach number computations are considered, a discretization scheme for incompressible flow equations has been extended to the isotropic compressible case. Hence, the approach is using primitive variables for the discretization of the Euler equations. For very low Mach numbers the resulting algebraic system of equations becomes stiff, which makes further demands on the smoother used in the multigrid cycle. Preliminary results will be shown.

Institution web page: http://www.uni-heidelberg.de/

May 8, 2001

Numerical Simulation in Science and Engineering

Gabriel Wittum
University of Heidelberg

Abstract:

Numerical simulation has become one of the major topics in computational science. To promote modeling and simulation of complex problems, new strategies are needed allowing for the solution of large, complex model systems. Crucial issues for such strategies are reliability, efficiency, robustness, usability, and versatility.

After discussing the needs of large-scale simulation we point out basic simulation strategies such as adaptivity, parallelism, and multigrid solvers. These strategies are combined in our simulation system UG (for “Unstructured Grids”).

In the second part of the seminar we show the application of these strategies to the simulation of processes from fluid dynamics and subsurface modeling. In particular we will discuss the construction of the Filtering Algebraic Multigrid methods (FAMG) and show an application of it for the simulation of bioremediation of a chlorine-contaminated aquifer.

Institution web page: http://www.uni-heidelberg.de/

May 4, 2001

Recent Advances in Direct Numerical Simulations (DNS) of Turbulent Shear Flows Laden with Particles (droplets or bubbles)

Said Elghobashi
University of California, Irvine

Email: selghoba@uci.edu

Abstract:

Particle-laden turbulent flows are encountered in nature and numerous engineering applications. Sand and dust storms near the surfaces of Earth and Mars particulate pollutants in the atmosphere and droplet sprays in combustors are some examples. Recent advances in parallel computer technology allowed the direct numerical solution (DNS) of the governing equations of these flows with minimum approximation. There are at present two formulations of DNS of particle-laden flows: Eulerian-Lagrangian (or trajectory) and Eulerian-Eulerian (or two-fluid). When the volume fraction of the dispersed particles is of the order of 0.0001 or greater, the particles modify the turbulence structure; this phenomenon is known as two-way coupling.

The objective of this seminar is to present an overview of the above two approaches and DNS results for some particle-laden and bubble-laden turbulent flows.

Speaker's web page: http://kolmog.eng.uci.edu/

Research web page: http://mae.eng.uci.edu/

Institution web page: http://www.uci.edu/

April 27, 2001

Global Data Distribution in Software Distributed Shared Memory Systems

David Lowenthal
University of Georgia

Email: dkl@cs.uga.edu

Abstract:

Distributing data is one of the key problems in implementing efficient distributed-memory parallel programs. The problem becomes more difficult in programs where data redistribution between computational phases is considered. The global data distribution problem is to find the optimal distribution in multi-phase parallel programs. We are investigating this problem in the context of a software distributed shared memory (SDSM) system.

In this talk, we first describe SUIF-Adapt, which is an integrated compiler/run-time system for finding efficient global data distributions, and then we discuss performance of SUIF-Adapt programs. We next describe how SUIF-Adapt obtains accurate redistribution time estimates. Obtaining these estimates is difficult because it depends on access patterns, page locations, and the SDSM consistency protocol. However, having accurate redistribution times is critical to choosing an efficient distribution. We discuss our integrated compiler/run-time method for finding accurate time estimates and describe optimizations that improve the performance of our algorithm.

Speaker's web page: http://www.cs.uga.edu/~dkl/

Institution web page: http://www.uga.edu/

April 17, 2001

A Survey of Some Recent Results on Two-Level Finite Element Preconditioning Methods

Owe Axelsson
University of Nijmegen

Email: axelsson@sci.kun.nl

Abstract:

Two-level methods for conforming finite element approximations of self-adjoint scalar or systems of elliptic equations are considered. Although this topic was first treated some twenty years ago, there are some recent results and there remain some open questions. The matrix that arises is naturally partitioned in a two-by-two block form $A=[Aij]$, $i,j=1,2$. The simplest but still powerful preconditioner is the block diagonal part, i.e., B = blockdiag (A11, A22). The corresponding spectral condition number is $(1 + \tilde{a})/(1 - \tilde{a})$, where $\tilde{a}$, $0 < \tilde{a} < 1$ is the constant in the strengthened CBS (Cauchy–Buryakowski–Schwarz) inequality. Equivalently, $\tilde{a}$ is the cosine of the angle between the subspaces spanned by the coarse mesh basis functions and the basis functions for the added node points, respectively.

The method can be extended by recursion to a multilevel method by approximating the coarse mesh matrix (A22) in a similar way. In this way, the resulting condition number can be stabilized using certain matrix polynomial approximations of Schur complement matrices.

Two major partly open questions arise: (1) the computation of the constant $\tilde{a}$ for three-dimensional bodies and the dependence of this constant on the diffusion coefficients; (2) preconditioning of the block matrix corresponding to the added degrees of freedom on each level to achieve a condition number which is bounded uniformly in the problem parameter and the level number.

Examples of some new recent results on each topic will be presented and some open questions will be discussed.

Speaker's web page:

http://www-math.sci.kun.nl/math/medewerkers/Axelsson.html

Institution web page: http://www.kun.nl/

April 9, 2001

A Review of Computational Modeling in the MHD Laboratory at STRELA

Oleg Diyankov
STRELA Open Computer Center

Abstract:

We present three topics in computational physics. The first is a brief description of the MHD model used for "2.5D" modeling. The mathematical approach to the solution of the 2.5D MHD equations will be discussed. Application of our code to Z-Pinch implosion problems and the treatment of magnetic field influence on laser-produced plasma motion will be presented. The second part of the talk is devoted to the description of a 2D free Lagrangian code. The mathematical approach to the solution of the Lagrangian gas dynamic equations on a moving grid, consisting of Voronoi cells, will be discussed. Example problems solved with the use of the code will be presented. The third part concerns the models and ideas that we would like to implement in the frame of an ISTC project, "Computer Modeling of Ceramic Pressing—the Creation of Ultra-high-hardness Materials." Our ideas on high order difference schemes for irregular grids will be discussed.

April 9, 2001

Numerical Simulations of Flows with Two Viscous Liquids

Yuriko Renardy

Virginia Polytechnic Institute and State University

Email: renardyy@math.vt.edu

Abstract:

We use a volume-of-fluid method to track the deformation and breakup of the interface separating two viscous liquids. The original code is Stephane Zaleski's SURFER. To this we have added a semi-implicit Stokes solver to handle low Reynolds numbers, parallelization to handle large 3D simulations, and new physics such as a model for a yield stress liquid, a model for surfactants, and moving contact lines.

Two simulations will be described. One is the axisymmetric simulation of bamboo waves for core-annular flow of oil and water. The second is rupturing of a viscous drop under simple shear including the effect of inertia. Videos will be shown.

Speaker's web page: http://www.math.vt.edu/people/renardyy/

Institution web page: http://www.vt.edu/

April 9, 2001

Textbook Multigrid Efficiency in Solution of Reynolds-Averaged Navier–Stokes Equations

Boris Diskin
ICASE, NASA Langley

Email: bdiskin@icase.edu

Abstract:

Full multigrid (FMG) algorithms are the fastest solvers for elliptic problems. These algorithms can solve a general discretized elliptic problem to the discretization accuracy in a computational work that is a small (less than 10) multiple of the operation count in one target-grid residual evaluation. Such efficiency is known as textbook multigrid efficiency (TME). The difficulties associated with extending TME for solution of the Reynolds-averaged Navier–Stokes (RANS) equations relate to the fact that the RANS equations comprise a system of coupled nonlinear equations that is not, even for subsonic Mach numbers, fully elliptic, but contains hyperbolic parts. TME for the RANS simulations can be achieved if the different factors contributing to the system can be separated and treated optimally, e.g., by multigrid for elliptic factors and by downstream marching for hyperbolic factors. One means of separating the factors is the distributed relaxation approach proposed by A. Brandt. Earlier demonstrations of TME solvers with distributed relaxation have been performed for incompressible free-stream inviscid and viscous flows without boundary layers.

In this talk, a general framework for achieving TME in solution of the high-Reynolds-number Navier–Stokes equations will be outlined. TME distributed-relaxation solvers will be demonstrated for viscous incompressible and subsonic compressible flows with boundary layers and for inviscid compressible transonic flows.

Speaker's web page: http://www.icase.edu/~bdiskin/

Institution web page: http://www.icase.edu/

March 20, 2001

Boundary Conditions and Estimates for the Stokes Equations on Staggered Grids

Bertil Gustafsson
Stanford University

Email: bertil@sccm.stanford.edu

Abstract:

We consider the steady state Stokes equations, describing low speed flow, and derive estimates of the solution for various types of boundary conditions. When using the equations in their original form with the divergence condition as part of the system, there is a necessary restriction on the boundary data. We formulate the boundary conditions in a new way, such that this problem is eliminated, and the boundary value problem becomes nonsingular and well conditioned. By using a difference approximation on a staggered grid, we are able to derive a nonsingular approximation in a direct way, and such that we get the same type of estimates as for the continuous case. Numerical experiments confirm the theoretical results.

In a recent thesis by Jonas Nilsson, the method is generalized to the time-dependent problem, to curvilinear grids, and to the incompressible Navier–Stokes equations. We will present a brief survey of this work as well.

Institution web page: http://www.stanford.edu/

March 14, 2001

Fast High-Order Methods in CEM

Oscar Bruno
California Institute of Technology

Email: bruno@acm.caltech.edu

Abstract:

I will discuss a class of fast, high-order solvers for problems of electromagnetic and acoustic scattering. Our methods include approaches based on perturbation of boundaries and frequencies, as well as novel fast, high-order integral techniques. I will present numerical results for a variety of computational problems containing high frequencies, complex scatterers, and geometric singularities such as edges and corners. (Collaborators: M. Caponi, M. Hyde, L. Kunyansky, F. Reitich, A. Sei)

Speaker's web page: http://www.acm.caltech.edu/%7Ebruno/

Research web page: http://www.acm.caltech.edu/

Institution web page: http://www.caltech.edu/

February 28, 2001

Optimal Preconditioned Eigensolvers for Very Large Symmetric Eigenproblems

Andrew Knyazev
University of Colorado, Denver

Email: andrew.knyazev@cudenver.edu

Abstract:

Many applied and engineering problems, e.g., in structural mechanics, design of nuclear reactors, ocean modeling, and quantum chemistry, lead—after simplifications and an appropriate approximation of original partial differential equations—to extremely large and ill-conditioned linear systems with symmetric positive definite matrices of coefficients and similar symmetric eigenvalue problems. The preconditioned conjugate gradient method became the standard solver for such linear systems. Our ultimate goal is to develop an analogous optimal method for eigenproblems. Ideally, we want to be able to compute an eigenvector of interest at the same cost as that of solving a linear system of equations, using the same preconditioner. That would allow, in particular, a simple adaptation for eigenproblems of available domain decomposition-based and multigrid preconditioners for linear systems.

Searching for the optimal eigensolver, we describe the Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method for symmetric eigenvalue problems, based on a local Rayleigh–Ritz optimization of a three-term recurrence. LOBPCG can be viewed as a nonlinear conjugate gradient method of minimization of the Rayleigh quotient, which takes advantage of the optimality of the Rayleigh–Ritz procedure. Numerical results establish that our LOBPCG method is practically as efficient as the best possible algorithm on the whole class of preconditioned eigensolvers. We discuss several competitors, namely, some inner-outer iterative preconditioned eigensolvers. Direct numerical comparisons with one of them, the inexact Jacobi–Davidson method, show that our LOBPCG method is more robust and converges almost two-times faster. Finally, we show numerically that the LOBPCG method is robust with respect to variable preconditioning. A MATLAB code of the LOBPCG method and the Preconditioned Eigensolvers Benchmarking are available at http://www-math.cudenver.edu/~aknyazev/software/CG/.

Speaker's web page: http://www-math.cudenver.edu/~aknyazev/

Institution web page: http://www.cudenver.edu/

February 16, 2001

Combinatorial Morse Theory and Persistence, with Applications to Distinguishing Noise from Feature

John Harer
Duke University

Email: John.Harer@Duke.edu

Abstract:

This talk will describe a combinatorial approach to analyzing data sets in two and three dimensions. Based on a discrete version of Morse Theory, the approach is first to build a combinatorial framework for the data, and then to work with that framework numerically. This approach is useful in many cases in controlling error accumulation. The talk will also discuss the theory of persistence as a mechanism for distinguishing noise from feature in data sets and will outline procedures for algorithm development and data handling that fit this approach.

Speaker's web page: http://www2.math.duke.edu/faculty/harer/

Institution web page: http://www.duke.edu/

February 9, 2001

Accounting for Stability: Accurately Estimating the Error of Numerical Solutions of Differential Equations

Donald Estep
Colorado State University

Email: estep@math.colostate.edu

Abstract:

Accurately estimating the error of numerical solutions of differential equations remains an important scientific problem. Recently, we have made significant progress using a new approach, which at its heart is computational rather than analytical. This approach is based on a variational a posteriori error analysis that takes into account both the local introduction of discretization error and the accumulation of errors. I will explain the ingredients of this theory and illustrate its application using a variety of examples.

Speaker's web page: http://www.math.colostate.edu/~estep/

Institution web page: http://www.colostate.edu/

January 5, 2001

Exact Analysis of the Cache Behavior of Nested Loops

Erin Parker

University of North Carolina, Chapel Hill

Email: parker@cs.unc.edu

Abstract:

We use Presburger arithmetic to exactly model the behavior of loop nests executing in a memory hierarchy. Our formulas can be simplified efficiently to count various kinds of cache misses and to determine the state of the cache at the end of the loop nest. Our model is powerful enough to handle imperfect loop nests and various flavors of nonlinear array layouts based on bit interleaving of array indices. We also indicate how to handle the modest levels of associativity found in current data caches, and how to perform a certain amount of symbolic analysis of cache behavior. The complexity of the formulas relates to the static structure of the loop nest rather than to its dynamic trip count, allowing our model to gain efficiency in counting cache misses by exploiting repetitive patterns of cache behavior. Validation against cache simulation confirms the exactness of our formulation. Our method can serve as the basis for a static performance predictor to guide program and data transformations to improve performance.

Institution web page: http://www.unc.edu/

January 4, 2001

A Compiler-Based Approach to Specializing Software Libraries

Calvin Lin
University of Texas, Austin

Email: lin@cs.utexas.edu

Abstract:

Software libraries are a popular means of reusing code. Unfortunately, the static black-box nature of libraries can thwart compiler optimizations, limit code reuse, and threaten portability. This talk introduces a new compiler-based approach to optimizing both software libraries and the application programs that use them. The key is a simple annotation language that describes semantic information about libraries. The result is a compiler that can analyze and transform library operations in the same way that language primitives can be analyzed and transformed. Experiments with the PLAPACK parallel linear algebra library show that this technique can yield significant performance improvements, even for a library that has been carefully designed to provide good performance. One goal of this talk is to identify potential users of our system and to meet with programmers who use libraries written in C or Java.

Speaker's web page: http://www.cs.utexas.edu/users/lin/

Research web page: http://www.cs.utexas.edu/users/lin/projects.html

Institution web page: http://www.utexas.edu/

November 30, 2000

A Particle Method and Adaptive Treecode for Vortex Sheet Motion in 3D Flow

Robert Krasny
University of Michigan

Email: krasny@math.lsa.umich.edu

Abstract:

A particle method is presented for computing vortex sheet motion in three-dimensional flow. The particles representing the sheet are advected using the Rosenhead–Moore form of regularized Biot–Savart kernel and new particles are inserted during the simulation to maintain resolution as the sheet rolls up. The particle velocities are computed using a treecode based on Taylor approximation in Cartesian coordinates. The necessary coefficients are computed using a recurrence relation, and several adaptive techniques are employed to gain efficiency, including nonuniform rectangular cells, variable order approximation, and a run-time choice between approximation and direct summation. Tests show that the treecode is significantly faster than direct summation for systems having a large number of particles. The algorithm is applied to simulate flows in which a vortex sheet rolls up to form a vortex ring. Two examples are presented: (1) azimuthal waves on a vortex ring, and (2) merger of two vortex rings.

(This is joint work with Keith Lindsay, NCAR.)

Speaker's web page: http://www.math.lsa.umich.edu/~krasny/

Institution web page: http://www.umich.edu/

November 28, 2000

Checking System Rules Using System-Specific, Programmer-Written Compiler Extensions

Dawson Engler
Stanford University

Email: engler@csl.stanford.edu

Abstract:

Systems software such as OS kernels, embedded systems, and libraries must obey many rules for both correctness and performance. Common examples include, "Accesses to variable A must be guarded by lock B," "System calls must check user pointers for validity before using them," and "Disabled interrupts must be re-enabled." Unfortunately, adherence to these rules is largely unchecked.

This talk shows system implementors can use "meta-level compilation" (MC) to write simple, system-specific compiler extensions that automatically check their code for rule violations. By combining domain-specific knowledge with the automatic machinery of compilers, MC brings the benefits of language-level checking and optimizing to the higher, "meta" level of the systems implemented in these languages. This talk presents results of applying MC to four real systems: Linux, OpenBSD, the Xok exokernel, and the FLASH machine's embedded software. MC extensions found over 600 errors in these systems and led to numerous kernel patches. Most extensions were less than a hundred lines of code and written by implementors who had a limited understanding of the systems checked.

Speaker's web page: http://www.stanford.edu/~engler/

Institution web page: http://www.stanford.edu/

November 27, 2000

About Mesh Generation and Adaptive Multigrid Methods

Joachim Schoeberl
Texas A&M University

Email: joachim@math.tamu.edu

Abstract:

Numerical simulation of complicated processes needs the interaction of different scientific disciplines. Here, the value of software design must not be underestimated. Especially, adaptive multilevel algorithms on complex 3D geometries require a careful design as well as a certain amount of manpower.

The first step in the simulation is the description of the geometry. From the geometry database, the mesh generator automatically builds the initial mesh. To gain optimal advantage of adaptive and hierarchical algorithms, the initial mesh shall be as coarse as possible to resolve the geometry. The geometry database and the mesh database, together with the problem description, are fed into the actual simulation program. Based on the approximate solution, a posteriori error indicators mark elements for refinement or coarsening. Hierarchical mesh refinement algorithms have to perform the changes and maintain information needed for multilevel solvers. It is important to have access to the geometry to adapt the mesh to the boundary and edges.

We will present available tools for the generation of the initial mesh and for performing the necessary mesh refinement. The underlying geometry description is based on a Constructive Solid Geometry (CSG) tree and includes now various types of spline curves and surfaces as well. The mesh generator NETGEN performs the steps from scanning the geometry until generation of the volume mesh. Small details in the geometry are detected and the mesh is automatically graded. This mesh can be used by any simulation program, which also sets the refinement flags. NETGEN then performs the necessary updates and returns the mesh together with the hierarchical information to the simulation code. There are different realizations of the interface, from file transfer to direct data access.

We will present applications, where NETGEN is used together with our fem code for the simulation of mechanical, electromagnetical, as well as coupled field problems.

Speaker's web page: http://www.math.tamu.edu/~joachim/

Institution web page: http://www.tamu.edu/

November 16, 2000

Janus: A Library for Adaptive Unstructured Scientific Applications

Peter Gottschling
GMD First

Email: pg@first.gmd.de

Abstract:

Janus is a conceptual framework and a C++ template library for irregular and regular scientific applications. It provides (potentially distributed) data structures to represent spatial structures and numerical data of these applications. Janus is implemented on top of the C++ Standard Template Library and uses MPI as default parallel platform. The components of Janus cooperate well with existing linear solvers and mesh partitioning packages, such as PETSc and ParMetis.

Janus rests on the observation that there occur two essential kinds of objects in scientific applications. The first kind is referred to as "spatial structures" such as (rectangular) grids, meshes, or graphs. The second kind is simulation data that are associated with the spatial structures. Typical examples are grid functions (vectors), element matrices on finite element meshes, and other (sparse) matrices. An important observation is that the spatial structures are conceptually more stable than the associated data. Janus takes advantage of this stability by putting as much computational effort as possible (load balancing, setting local index schemes, etc.) in the construction of spatial structures. On the other hand, the access to the spatial structures and the usage of associated data are fast.

The conceptual framework of Janus is designed using the paradigm of generic programming that has been successfully applied in the C++ standard library and other libraries such as the Matrix Template Library (MTL) or the POOMA framework. Using templates enables a more problem-oriented programming style, i.e., domains of arbitrary types, like triangles or hexahedra, can be defined. Furthermore, relations are defined between arbitrary domains, not between index sets.

For more information about Janus visit: http://www.first.gmd.de/promise/

Research web page: http://www.first.gmd.de/~jens/janus/

Institution web page: http://www.first.gmd.de/Welcome.ORIG.html

November 2, 2000

Virtual Global Grids for Adaptive Methods in Space and Time

Ulrich Ruede
University of Erlangen-Nuernberg

Email: ruede@informatik.uni-erlangen.de

Abstract:

Many advanced methods use adaptive techniques. In this talk we will present a time-space adaptive discretization and iterative solution technique. The method is based on the fully adaptive multigrid method (FAMe) and virtual global grids, which have been proposed for the adaptive solution of elliptic equations. Here, adaptive discretization and iterative solution are integrated in one process. The approach can also be extended to efficiently exploit locality as needed, e.g., for an efficient parallelization or cache-aware algorithms.

In the time-dependent case, we propose not to use global time steps, but to treat time with a multilevel approach, similar to time parallel multigrid methods. Thus, the iterative solver iterates on all time levels simultaneously. We will discuss how this can be extended to use adaptive meshes in space and time. This method has been implemented and tested for prototype model problems.

Speaker's web page: http://www10.informatik.uni-erlangen.de/~ruede/

Institution web page: http://www.uni-erlangen.de/

October 26, 2000

Algebraic Multilevel Preconditioning Based on Approximate Cyclic Reduction

Johannes Kraus
University of Leoben

Email: jkraus@poptop.llnl.gov

Abstract:

We consider preconditioners for large sparse matrices arising from discretization of partial differential equations of predominant elliptic type. The proposed purely algebraic multilevel method is based on approximate cyclic reduction. Within an incomplete LU, decomposition process spanning trees of matrix graphs are constructed that rest on a local optimization principle. A red-black coloring of these subgraphs yields the partitioning into fine- and coarse-grid variables and is also utilized to determine appropriate approximations of the Schur complements (the coarse-grid operators) on different coarse levels. This idea is combined with algebraic multilevel iterations, V- or W-cycle, in the preconditioning step of some Krylov subspace method. Another variant defines a nonlinear algebraic multilevel iteration of W-cycle type based on inner GCG iterations. This results in a parameter-free method that shows similar convergence properties and optimal order of computational complexity for various model problems.

Institution web page: http://www.unileoben.ac.at/

October 24, 2000

Semi-Unstructured Grids in 3D

Christoph Pflaum
University of Wurzburg

Email: pflaum1@llnl.gov

Abstract:

It is well known that structured grids have several advantages in comparison to unstructured grids (e.g., high performance, low storage requirement, local refinement, and construction of coarse grids). The disadvantage of pure structured grids is the difficulty of constructing an automatic grid generator for structured grids on general domains. To avoid this problem one can use semi-unstructured grids. Such grids consist of a maximal structured grid in the interior of the domain and a subdivision of boundary cells by tetrahedra. It will be shown that the number of types of boundary cells in such a construction can be as few as twelve, if the domain does not contain slits. Furthermore, each boundary cell can be subdivided by tetrahedra such that the maximum interior angle is 163 degrees in case of a general domain and 145 degrees in case of a convex domain. Theoretical and numerical results are presented.

October 19, 2000

Osiris Analysis: A Diagnostic and Visualization Package for Large Data Sets from Particle in Cell Codes

Ricardo Fonseca
University of California, Los Angeles

Email: zamb@physics.ucla.edu

E.S. Dodd, R.G. Hemker, S. Lee, C. Ren, F. Tsung, L.O. Silva, W.B. Mori

Abstract:

Visualization is an essential part of any numerical laboratory. The data being produced by present day simulations is both large and complex, and it requires an appropriate set of tools for the adequate understanding of the results obtained. We present here the Osiris Analysis package, developed for the visualization and analysis of scientific data from Particle-In-Cell (PIC) simulations. This package consists of a wide set of visualization routines for scalar and vector data up to three dimensions in size, and it can also be used for any type of grid-deposited data and general particle data diagnostics. The analysis package includes spectral analysis routines, envelope and centroid analysis, *n*-dimensional local peak analysis, and a Poisson solver. This package was developed using the Interactive Data Language (IDL) software package and includes a set of interfacing routines based on the HDF file standard. We will discuss applications from Osiris 3D simulations for the Weibel Instability and Intense Laser/Beam propagation in plasmas.

Speaker's web page: http://exodus.physics.ucla.edu/

Institution web page: http://www.ucla.edu/

October 18, 2000

Algorithm Development for the Stanford University ASCI TFLO Solver

Antony Jameson (with Juan Alonso)
Stanford University

Email: jameson@baboon.stanford.edu

Abstract:

In this talk we describe the efforts, objectives, and methods of the ASCI turbomachinery simulations group at Stanford University. Our group is in charge of the development of efficient solvers for the computation of large-scale unsteady flow simulations in compressors and turbines with multiple stages using a Reynolds Averaged Navier–Stokes model. We begin by presenting a description of our objectives and relating some of the details of the flow solver, TFLO, which is being used for this project. In particular, the dual-time stepping, multigrid-driven algorithm is explained in some detail together with its performance, potential, and shortcomings. We will also present the results of a variety of unsteady large-scale test cases that we are currently conducting on ASCI platforms. Finally, we discuss two alternative approaches that can help speed up the computation of unsteady flows. These approaches are based on a modified ADI/dual-time scheme and on a harmonic balance technique.

Research web page: http://cits.stanford.edu/

Institution web page: http://www.stanford.edu/

October 16, 2000

Diffpack Object-Oriented Software Solutions for PDEs

Are Magnus Bruaset
Numerical Objects AS

Email: amb@nobjects.com

Abstract:

There is rapidly growing interest in the use of modern software technology and object-oriented principles in the context of scientific computing. The C++ software framework Diffpack is the result of more than ten years of dedicated research and development in this direction, with a special emphasis on numerical solution of partial differential equations.

Diffpack is today a commercial product from Numerical Objects in Oslo, Norway. The software is under continuous development in close collaboration with the Scientific Computing group at the University of Oslo.

The talk will focus on the underlying principles of Diffpack and will show several application examples addressing academic as well as industrial problems. These applications are taken from a wide range of areas, including traditional engineering disciplines like structure mechanics, computational fluid dynamics and porous media flow, as well as computational medicine and computational finance.

We will present recent advances in the software, with special emphasis on the Diffpack Parallel Computing Toolboxes. This functionality, which will ship early next year, allows a very easy transition from a sequential to a parallel Diffpack simulator. We will present results from using this software approach in a wide range of applications on a Linux-based cluster of workstations, as well as in traditional parallel environments.

Institution web page: http://www.nobjects.com/

October 5, 2000

Coupled Flows with a Focus on the Eye

Victor Barocas
University of Minnesota, Twin Cities

Email: baroc001@tc.umn.edu

Abstract:

Coupled flows, i.e., those that involve flow of a fluid past a deformable solid, present a number of challenges, but they are extremely common and thus of considerable importance. The greatest challenge associated with the coupled problem is that the domain for the fluid flow problem is not known until one has solved the solid deformation problem, which is in turn coupled to the fluid problem by drag-induced deformation. We use a pseudo-solid approach in which an imaginary compressible solid (the pseudo-solid) is overlaid onto the fluid domain, and the motion of the domain is tracked by monitoring the motion of the pseudo-solid. We have used the pseudo-solid approach to solve the coupled flow problem associated with the passive deformation of the iris (the washer-shaped colored part of the eye) in response to flow of aqueous humor (the clear fluid that bathes the iris). I will present results for the axisymmetric form of the problem with application to healthy and diseased eyes, and I will also discuss our computational time issues (especially for 3D) and our work with T. Manteuffel and S. McCormick of the University of Colorado to develop faster solution methods.

Research web page: http://www.bme.umn.edu/

Institution web page: http://www1.umn.edu/twincities/